Statistical Methods for Hazards and Health

by Yvonne M. M. Bishop*

The objective of this article is to document the need for further development of statistical methodology, training of more statisticians and improved communication between statisticians and the many other disciplines engaged in environmental research. Discussion of adequacy of the current statistical methodology requires the use of examples, which will hopefully not be offensive to the authors. Reference is made to recent developments and areas of unsolved problems delineated in three broad areas: enumeration data and adjusted rates; time series; and multiple regression.

A brief outline of the ideas behind current methods of analyzing discrete data is followed by a demonstration of their utility using an example of the effects of exposure, sex, and education on bronchitis rates.

Examples are listed of the ubiquity of the time component when relating pollution effects to each other and to health effects. An artificial example is used to emphasize the effects of time-dependent autocorrelations, trends, and cycles. References are given to a variety of new developments in time-series analysis.

Discussion of the pitfalls in multiple regression analysis, and possible alternative approaches is largely based on two recent reviews and includes references to recent developments of robust techniques.

Introduction

Dramatic episodes of fog or smog accompanied by notably increased mortality and morbidity have convinced us that polluted air affects health (*1–3*). Now we must determine more precisely how much pollution and what type of pollution causes disability. Both the exposure variable "air quality" and the outcome variable "health effects" are hard to define and measure. Much discussion centers on the reliability and validity of specific measures; increasingly, attention is being paid to numerous ancillary factors or covariates that influence postulated relationships. All these issues are of crucial importance in designing good studies and point to the need for interdisciplinary input when studies are being designed. If a study is poorly designed no amount of subsequent statistical legerdemain will produce meaningful results. Conversely, even the best designed studies can lead to misleading conclusions if the data are inadequately analyzed. We need both good design and good analysis.

This paper addresses only the issue of data analysis and ignores study design, except insofar as improvements of analytic techniques will reflect on

design requirements. As the need for better methodology cannot be appreciated unless the deficiencies of the present state-of-the-art are considered, examples will be given where the information obtained from the available data is not optimum. Examples for this purpose have been taken from a Chess monograph (*4*). In some instances the state of the art has improved since this work was done; in other areas many deficiencies still exist. The purpose of using these examples is not to criticize but to demonstrate the importance of improving our analytic techniques.

The introductory overview to the Chess monograph cites two statistical methodologies, general linear regression for quantitative variables and general linear models for categorical responses (*4–6*). The similarity of the two methods is stressed. Below we show how the emphasis on this similarity has led the authors to report their analyses of categorical models inappropriately and generally inadequately exploit the strengths of the analytic technique. We discuss the problems of time series and why linear regression techniques are inappropriate for their analysis. Some of the modern advances in fitting linear and nonlinear models to quantitative variables are mentioned briefly. We conclude that the 1970 task force recommendations should be stressed once again.

*Harvard School of Public Health, Boston, Massachusetts 02115.

Enumeration Data and Adjusted Rates

What Is a Log-Linear Model?

In recent years there has been much development in the handling of discrete data that have many categorical variables. Most authors agree that the interactions between the variables can best be determined by fitting models that are linear in the logarithmic scale.

Suppose we are interested in the effect of the three variables sex, age, and exposure area on the prevalence of bronchitis. The most complex model states that each of the three variables has a proportional effect on the bronchitis rate, and that each pair of variables may modify the effect of the other, and indeed that all three variables may have a joint effect. This is equivalent to saying that the effect of age on the bronchitis rate is not the same for each sex, and that the magnitude of this interaction varies between exposure areas. We say that this model includes the four-factor interaction bronchitis–age–sex–area. At the other extreme, the simplest model states that the bronchitis rate is constant for every sex–age–area combination. Between the most complex and the simplest model we can choose from a large variety of intermediate models, each postulating different combinations of simple proportional main effects and interaction effects. Each main or interaction effect is represented by a term in the log-linear model. Analysis consists of determining which intermediate model fits the data well and is not appreciably improved by adding more terms.

How Do We Choose a Model?

Although most authors are agreed upon the general utility of the log-linear model approach, there is some disagreement over the methods of obtaining estimates under a specific model and determining how well these estimates fit the observed data. Most of the proposed methods such as maximum likelihood, least squares, or minimum chi-square usually yield comparable if not identical estimates, and the probability levels associated with the goodness-of-fit statistics are in general very close. Thus although we can chose from a variety of techniques for fitting models to a particular data set, the final selection of a suitable model is not dependent on the choice of technique. Further discussion of comparisons between techniques has been given elsewhere (7, 8).

A well-fitting model is selected by a process of trial and error, and it includes those main effects and interactions which are large. The main effects and interactions that do not improve the goodness-of-fit are discarded. We often declare that the effects that are included are "significant" and those that are discarded are "not significant." Indeed, we may finish up with a table resembling an analysis of variance table. Such a table will list effects of importance, and given an indication of how the overall goodness-of-fit would be changed if each effect is excluded from the model. The degrees of freedom associated with these measure-of-fit statistics are determined from the number of categories in the relevant variables. The most commonly used measures are asymptotically distributed according to the chi-square distribution and so the probability of observing a value as large or larger than value tabulated may be readily obtained.

How Does This Help Us?

Fitting models may be helpful in two ways: (a) we can determine which effects are of importance, and (b) we can use the fitted estimates obtained under the model in order to obtain meaningful summary statistics. In our example above, meaningful summary statistics might be bronchitis rates for each exposure area adjusted for differences in the sex and age distributions in the areas.

The models can be extended to include many variables. As an example of the type of situation where they are of value we include Tables 1-3 which are taken from the Rocky Mountain studies (4). Inspection of the first Tables 1 and 2 indicates that we have the following five variables: bronchitis, two categories, yes or no; sex, two categories; education, three categories; age, four categories; exposure area two categories.

Multiplying together the number of categories tells us that each person is distributed into one of 96 cells. It is difficult to interpret Table 3 because sufficient information on which model was fitted is not given. If we assume (a) that sex, education and age are related to bronchitis rates, (b) exposure area has no effect on bronchitis rates, (c) the numbers of persons in each sex–education–age category differs by exposure area, and (d) that no multifactor effects are present, then the model fitted would have the terms shown in Table 4, each with their associated degrees of freedom, one for each parameter.

Table 1. Smoking- and sex-specific prevalence rates (percent) for chronic bronchitis by education and age[a,b]

Category	Nonsmokers		Ex-smokers		Smokers	
	Mothers	Fathers	Mothers	Fathers	Mothers	Fathers
Education:						
<High school	2.06	3.23	5.83	4.81	14.50	21.18
High school	2.20	3.66	1.43	3.49	11.75	15.70
>High school	1.36	1.95	2.51	2.13	10.85	19.46
Age						
≤29	1.09	0.00	2.63	0.00	13.07	14.55
30–39	1.31	0.68	3.86	2.72	11.17	15.59
40–49	2.63	4.10	2.38	3.24	14.95	21.00
≥50	2.61	6.25	0.00	5.06	7.69	28.41

[a]Data from Chess monograph (4).
[b]Chronic bronchitis rates are equivalent to crude rates for symptom severities 6 and 7.

Table 2. Prevalence of chronic bronchitis in nonindustrially exposed parents: individual and pooled community rates (percent) by sex and smoking status[a]

Community	Nonsmokers		Ex-smokers		Smokers		Sex- and age-adjusted rates		
	Mothers	Fathers	Mothers	Fathers	Mothers	Fathers	Non-smokers	Ex-smokers	Smokers
Pooled low	1.08	1.25	3.12	1.45	11.78	17.05	1.08	2.46	14.00
Low I	1.36	1.10	3.05	0.00	8.72	12.44			
Low II	0.48	0.00	4.55	5.31	14.15	20.86			
Low III	1.06	4.88	0.00	0.00	13.68	20.00			
Pooled high	2.54	3.47	2.80	4.82	12.88	18.63	3.12	3.56	15.72
High I	3.56	4.90	1.79	4.72	13.83	18.40			
High II	1.50	2.00	3.92	4.95	11.75	18.85			

[a]Data from Chess monograph (4).

Table 3. Analysis of variance for health observations in smokers and nonsmokers, chronic bronchitis.[a,b]

Factor	Degrees of freedom	Smokers		Nonsmokers[a]	
		χ^2	Probability (p)	χ^2	Probability (p)
Sex	1	14.70	<0.0003	1.12	<0.3
Education	1	3.15	<0.07	0.48	<0.6
Age	1	8.61	<0.004	10.10	<0.002
Exposure	1	0.66	<0.5	8.73	<0.005
Fit of model	11	14.19	<0.2	13.85	<0.2

[a]Data from Chess monograph (4).
[b]Ex-smokers and lifetime nonsmokers were combined for this analysis to obtain a larger sample size.

Table 4.

Effects	Parameters fitted
Overall mean and adjustment for distribution of persons between areas	49
Bronchitis × sex effect	1
Bronchitis × education effect	2
Bronchitis × age effect	3
Total	55

If we fit a model with these 55 parameters to the 96 cells we have 96 − 55 = 41 degrees of freedom for assessing the goodness-of-fit of our model. By fitting models with each of the interaction effects removed in turn, we would have one, two, and three degrees of freedom associated with the differences in goodness-of-fit. Thus our table would not resemble Table 3 very closely. We would of course have only one degree of freedom for each effect if we reduced the number of categories in each variable

to two. With this reduction we would have 32 cells and be fitting 20 parameters, giving 12 degrees of freedom. The addition of the effect of exposure on bronchitis would bring us to 11 degrees of freedom as given in Table 3. This example has been cited laboriously to illustrate the importance of specifying which model was fitted.

There were further problems in understanding Table 3. Apparently two separate models were fitted, one to smokers and the other to nonsmokers. If we look at the first line of the table we see χ^2 values for sex and education are larger for smokers than for nonsmokers. We might suspect that smoking had a synergistic influence and enhanced the effects of age and education. Such a suspicion would be unjustified if the sample of smokers was larger than the sample of nonsmokers. We cannot make the assumption because χ^2 values increase with larger sample sizes, even when the interaction effect they reflect remains constant. We could readily evaluate the possibility of smoking affecting other interactions by the simple procedure of adding smoking as a sixth variable to the other five variables already in the model. Then we could determine the magnitude of possible three-factor effects—one relating smoking–sex–bronchitis and the other relating smoking–education–bronchitis.

If we turn to the second purpose of model fitting—to enable us to adjust rates for several underlying variables simultaneously—we find that this strength of the procedure has been ignored. All the rates given are either crude rates, or adjusted for at most two variables using crude specific rates.

What Improvements Are Needed?

In conclusion, the full strengths of the methodology were not used: (1) variables were reduced to two categories thus losing information, (2) smoking was not included as a variable, thus its effect cannot be assessed from the results given, (3) the particular model fitted could only be inferred, thus its goodness-of-fit statistics are of no value, (4) the fitted values were not used to compute adjusted rates. Some of the difficulties noted above stem from the attempt to present the results in a table format that resembles analysis of variance for continuous data. Although there are similarities in that models are being fitted, it is important to distinguish between the strengths of the different methodologies appropriate for different types of data (9). Thus the inadequacies were largely due to a lack of understanding of the methodology. This indicates a need for better training and communication.

Since 1970, further advances in technology have been made, notably methods for dealing with or-

dered categories (10–14) and methods for computing variances for certain types of estimates. There is still need for further development of methods suitable for a mixture of discrete and continuous variables.

Time Series

Why Do We Need to Look at Them?

The following are examples of situations where the relationships between two or more series of data collected over time are of current interest: (1) assessing the performance of a new pollution-measuring device compared with that of a standard device in the field; (2) determining whether adjacent stations monitoring the air in a city are giving comparable data or whether there are real differences in air quality in neighboring regions; (3) determining whether central monitoring stations give a true picture of individual exposure by comparing their readings with personal dosimeter readings; (4) relating fluctuations in indices of disease such as deaths, hospital visits or exacerbation of symptoms to measures of air quality; (5) assessing the extent to which different pollutants increase and decrease simultaneously or with a consistent lag between peaks; (6) prediction of the future levels of a given series so that the effects of intervention may be assessed.

Thus the relationship of various time series is central to relating environmental and health effects.

Why Is a Simple Correlation Not Informative?

In each of the situations cited above attempts have been made to use simple correlations as measures of the association between two time series. This approach can be criticized on several levels.

Range of Observations. If each serial measurement could be regarded as independent of all preceding measurements (which is usually untrue) and was taken from a normal distribution then correlation would be a reasonable approach. However when observing natural phenomenon the strength of the association will depend on the range of values that occurred during the observation period.

As an illustration, consider Figures 1a and 1b. In Figure 1a, two lines, marked A and B, are connecting a series of points. The points were obtained from a table of random normal deviates (15). Thus the points are independent observations from a normal distribution with mean of zero and variance
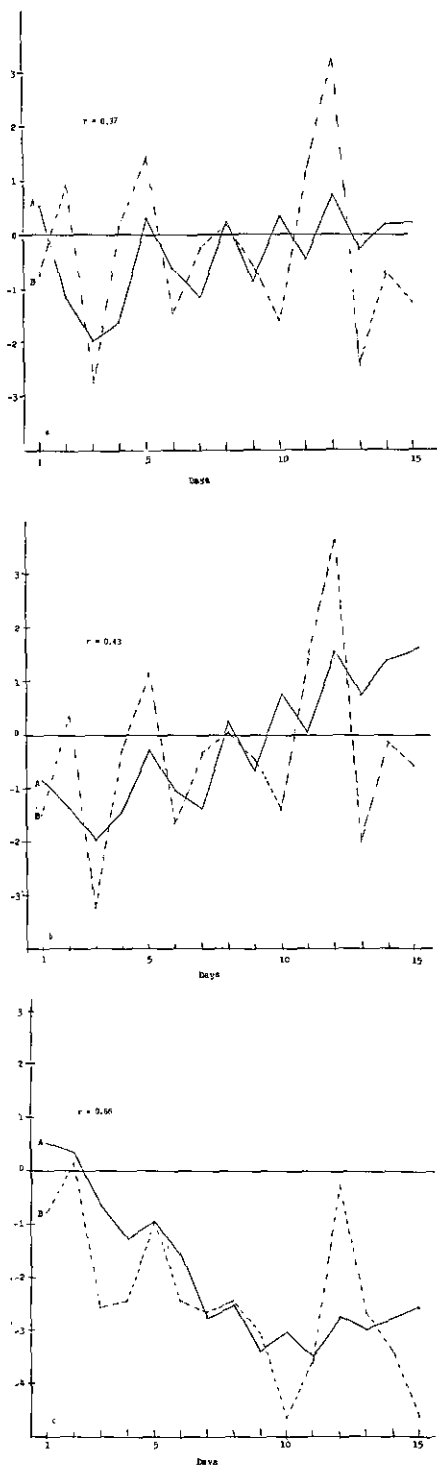
FIGURE 1. (a) Two series of independent normal deviates, r = 0.37; (b) same series as Fig. 1a with different trends added to each series, r = 0.43; (c) same series as Fig. 1a with autocorrelation within each series r = 0.66.

of one unit. Theoretically the two series of independent observations have a correlation of zero. By chance we have an observed value of $r = 0.37$. In Figure 1b we have introduced linear trends by adding to these random deviates a difference of 0.2 between successive measurements on line A, and differences of 0.1 for line B. The correlation we now compute is increased to $r = 0.43$. If we were to introduce steeper trends by adding larger constants, we would get larger values of r.

Clearly, in periods of relative stability of the underlying phenomena the values we obtain represent noise about the constant true value, as in Figure 1a, and the correlation between the two series will not differ significantly from zero. If we measure both phenomena during a period when both are subject to a seasonal trend, as in Figure 1b we will increase our apparent correlation. If we measure during a period when there is a period of stability and a period when both phenomena have a trend we will obtain an intermediate value for r. Before computing a correlation coefficient, it is necessary to consider whether the series have common large shifts, whether we need to distinguish short-term association from general seasonal trends and in fact to consider carefully the hypothetical model we are evaluating.

Successive Values Not Independent. Most of the time series data of interest cannot be regarded as independent observations as we did in the preceding section. We have only to consider a familiar measure such as minimum 24-hr temperature to appreciate that the possible values for a particular day fall within a range determined by knowing the time of year and can be defined even more closely by knowing the values for immediately preceding days. Thus the series are autocorrelated; the values for day t are related to those for day $t - 1$, and so on. This autocorrelation invalidates the use of regression or multiple regression techniques designed for independent observations. The effect of autocorrelation is shown in Figure 1c. The random value for each day in Figure 1a has been added to the value for the previous day, to provide a new series. We note that the new series looks smoother, as each day's values in a given series are related. The two series are, however, still unrelated to each other, except insofar as they have the same internal relationship. The observed correlation has however changed to $r = 0.66$.

Advances and Needs in Time Series Analysis

The foregoing simple examples illustrate some of the characteristics of time series that must be handled. Almost any series will exhibit noise and au-

tocorrelation, and most will have cyclic patterns of varying length.

Bloomfield (16, 17) has investigated the use of spectrum analysis as a tool for determining whether the aggravation of asthma symptoms are related to daily minimum temperature or to atmospheric SO_x levels. He explains: "The spectrum may be regarded as a decomposition of the variance of the data into components associated with different frequencies." Frequencies in this context means number of cycles per day; thus an annual effect would theoretically be at the frequency of 1/365 cycle per day, but in fact the smoothing of the data (which was a necessary preliminary step) spreads the effect over a wider band. Bloomfield also computes the coherence between series, which he explains as "the frequency-dependent measure of correlation between series." Thus he has a series of correlations that show the extent to which the cyclic patterns of the series correspond. He concludes, "the series are essentially unrelated at frequencies above 0.25 cycles per day, which correspond to a period of four days. However, at lower frequencies, which correspond to longer periods, there is substantial coherence. This is a warning that the impact of these two series on the health series may be complex and hard to disentangle."

He also investigates partial coherence, namely the frequency-dependent partial correlation between asthma and sulfur oxide after correction for the effect of minimum temperature. Throughout his paper he warns us about assumptions underlying the analysis, namely that the series are "stationary" in the sense that the covariances between time periods are constant throughout the series, and that the relationships between the variables are linear, and finally that the tentative conclusions reached may be reversed following subsequent analysis. Thus we conclude that this is a very promising approach but that care must be taken to recognize the importance of the underlying assumptions.

Stressing the limitations of a particular model is not intended to indicate that the approach is poor—rather it is to stress that analysis of time series is not simply a matter of running the data through a computer program. The situation is described by Box et al. (18): "The obtaining of sample estimates of the autocorrelation function and the spectrum are non-structural approaches, analogous to the representation of an empirical distribution function by a histogram . . . They provide a first step . . . pointing the way to some parametric model on which subsequent analyses will be based.

Box and other authors (19–21) have been developing such specific models for carbon monoxide in Los Angeles to study the effect of changes in methods of instrument calibration and the effect of various control measures.

The noise inherent in any system together with the limitations of the lengths of the series, usually requires that some form of smoothing is carried out during the analysis. Researchers at Princeton have been making rapid advances in development of these techniques and are conducting Monte Carlo simulations to evaluate different approaches. Thus again the research is in progress but much needs to be done before the relative advantages of different strategies are fully understood (22–24).

Multiple Regression

When Are Least-Squares Fits a Poor Choice?

Pitfalls in the interpretation of linear least-squares regression relating to two variables are well known; they include nonnormality of the distribution of variables, nonlinearity of the relationship between the variables, lack of independence between observations and the presence of outliers. When the number of variables increases so do the problems: the list must be enlarged to include multicollinearity of the variables, and it is no longer possible to detect these problems by simple plots of the data. Even when the problems are detected, the optimum method of analyzing data with one or more types of departure from the assumptions underlying least-squares regression is not readily apparent. Recent developments deal with both methods of detecting particular types of departure and with data-analysis in the presence of such departures. Increasingly these methods are being applied to analysis of environmental data but are apparently not well known to all investigators.

Directions of Current Development

In a recent review, Hocking, (25) suggests that "the role of the developers of regression methodology is to provide the less skilled user with techniques that are robust while easy to use and understand." Much effort has gone into the development of techniques that are "robust," or, in other words, are relatively insensitive to departures from the usual assumptions underlying least-

squares regression. Gnanadesikan et al. (26) have been particularly concerned with the detection of outliers. Andrews (27, 28) has re-analyzed data originally analyzed by Daniel and Woods (29), using newer techniques that he believes are resistant to a small number of gross outliers. He warns that his iterative technique is more expensive than least-squares but in addition to producing stable estimates it will detect outliers. Andrews reaches the same conclusions regarding this sample data set as Daniel and Wood, and this has led Hocking (25) to observe that these skilled analysts using repeated inspection of residual plots were in fact using a robust procedure. Diaconis (30) has applied resistant analysis of variance techniques to air pollution data. Brown et al. (31) observed reduction in mortality rates in two California counties and suggested that this might be a reflection of reduced air pollution consequent upon the 1974 fuel crisis. Diaconis was unable to find parallel reduction in CO or NO_2. Thus the question remains open whether the observed reduction in mortality was due to other causes, or to chance fluctuations, or to interactions among air pollutants that have not yet been investigated.

The problem of multicollinearity has been tackled by a variety of approaches. Schwing and McDonald (32) have compared least-squares and ridge regression, and have applied both ridge regression and a sign-restricted least-squares method to the analysis of the association between mortality rates, natural ionizing radiation, and some air pollutants. They show that the two later approaches yield comparable results that differ from those obtained by using least-squares (32, 33). The implications of order restrictions have also been investigated (34). In the conclusion of his review Hocking (25) states that "the multicollinearity problem seems to have been given too little attention in the statistics literature." He recommends that eigenvalues should always be inspected to determine possible redundancies, but that when near-singularities exist the method of handling them is not clear.

The problem of more complex relationships between variables has received much attention. In a recent review, Gallant (35) concentrates on methods of fitting nonlinear functions rather than on the detection of such functional relationships in the data. Other authors such as Anscombe (36), and Wilk (37), and Cleveland and Kleimer (38) have developed sophisticated plotting techniques for detection of characteristics of the data. Gnanadesikan and Kettenring (26) review many of these.

All of these endeavors point to the complexities that may be encountered in multivariate data. In view of these complexities, it is unlikely that a least-squares fit of a simple "hockey stick" function will prove to be an adequate method of determining "threshold" levels of pollutants as has been done (Fig. 2). This method may be useful in an experimental situation such as that described by McNeil (39), because other sources of variation are controlled. Certainly it is misleading to present point estimates obtained by this method without indicating their variability, and without reporting any attempt to investigate alternate models.
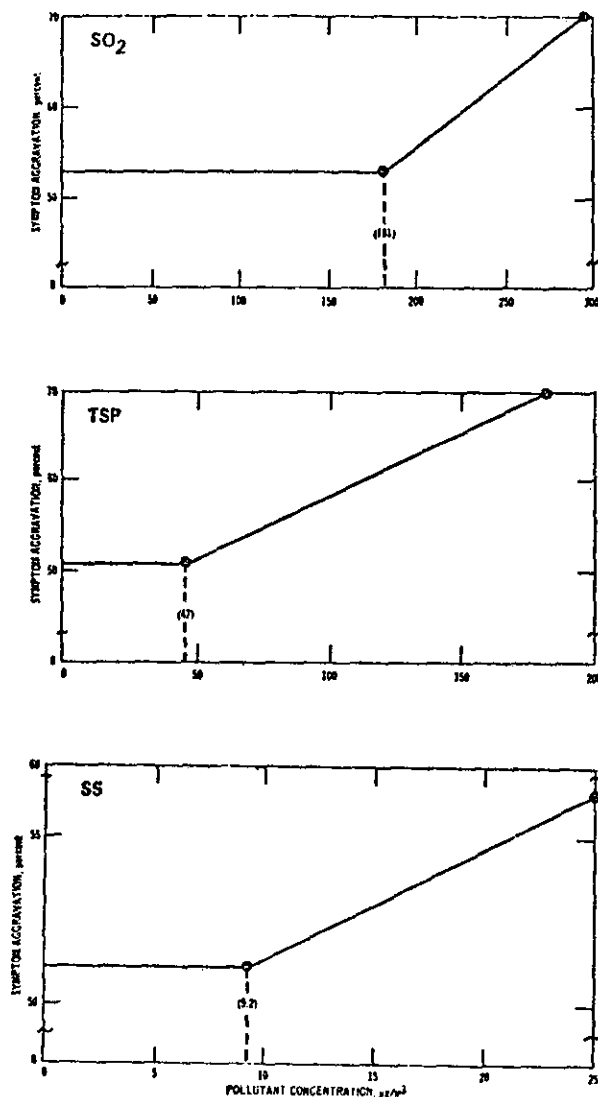

FIG. 2 Examples of the use of a hockey-stick function where no attempt is made to indicate reliability or to assess the interaction effects of different pollutants (4). The plots show temperature-specific threshold estimates for symptom aggravation by sulfur dioxide, total suspended particulates (TSP), and suspended sulfates (SS).

In the example reproduced in Figure 2 the effect of temperature was held constant, but three different pollutants were each treated separately with no attempt being made to consider how they would affect symptom aggravation when present in different combinations. Similar observations were made by the discussants of a paper by Nelson et al. (40).

Conclusions

The report of the task force on research planning in Environmental Health Sciences (41) recommended in 1970 that further development of efficient statistical techniques be undertaken. In at least three of the five areas of concern (contingency tables, time series, and multivariate methods), theoretical advances have been made. In some areas these advances have been well documented, in others progress has only reached the stage of verbal reporting and unpublished manuscripts. Much needs to be done, both in terms of development of theory and making readily accessible computer programs with adequate documentation for carrying out the techniques proposed.

In spite of this developmental activity, review of recent literature reveals relatively few instances where the newer techniques are being employed. Partly this is because the stage of development is such that they are not readily available, partly because of lack of communication. Thus the need for training recommended in 1970 still exists.

A satellite symposium was sponsored by IASPS on Statistical Aspects of pollution problems in 1971 (42). In the published report, Van Belle noted the dangers that "producers" of statistical analyses will base their product on arguments of dubious validity. He cites four areas: the first two were: (1) "The use of a linear regression model to approximate a cause–effect link is questionable" and (2) "The use of elasticity coefficients is misleading when the variables are measured in arbitrary units."

He also cautions about the indiscriminant accumulation of large bodies of data and on the tendency to place too much faith in "indices." These problems are still with us.

search Needs." Copies of the original material for this Background Document, as well as others prepared for the report can be secured from the National Technical Information Service, U.S. Department of Commerce, 5285 Port Royal Road, Springfield, Virginia 22161.

REFERENCES

1. Glaser, M., Greenberg, L., and Field, F. Mortality and morbidity during a period of high levels of air pollution, Nov. 23-25, 1966. Arch. Environ. Health 15: 684 (1967).
2. Schronk, H. H., et al. Air pollution in Donora, Pa.: Epidemiology of an unusual smog episode of October, 1948. P. H. Bulletin 306, Fed. Sci. Agency. Div. Ind. Hyg. PHS, U.S. Dept. HEW, Washington, D.C., 1949.
3. Scott, J. A. The London fog of December 1966. Med. Office. 109: 250 (1966).
4. Health Consequences of Sulfur Oxides: A Report from Chess, 1970-1971. U. S. Environmental Protection Agency, Research Triangle Park, N. C., 1974.
5. Graybill, F. An Introduction to Linear Statistical Models, Vol. 1. McGraw-Hill, New York, 1961.
6. Grizzle, J. E., Starmer, C. F., and Koch, G. G. Analysis of categorical data by linear models. Biometrics 25: 489 (1969).
7. Bishop, Y. M. M., Fienberg, S. E., and Holland, P. W. Discrete Multivariate Analysis: Theory and Practice. MIT Press, Cambridge, 1975.
8. Johnson, W. D., and Koch, G. G. A note on the weighted least-squares analysis of the Ries-Smith contingency table data. Technometrics 13: 438 (1971).
9. Wermuth, N. "Analogies between multiplicative models in contingency tables and covariance selection. Biometrics 32: 95 (1976).
10. Bock, R. D. Multivariate analysis of qualitative data. In: Multivariate Statistical Methods in Behavioral Research. McGraw-Hill, New York, 1975.
11. Clayton, D. G. Some Odds Ratio Statistics for the Analysis of Ordered Categorical Data. Biometrika 61: 525 (1974).
12. Haberman, S. J. Log-linear models for frequency tables with ordered classifications. Biometrics 30: 589 (1974).
13. Simon, G. Alternative Analyses for the singly-ordered contingency table. J. Amer. Statist. Assoc. 69: 971 (1974).
14. Williams, O. C., and Grizzle, J. E. Analysis of contingency tables having ordered response categories. J. Amer. Statist. Assoc. 67: 55 (1972).
15. Rand Corporation. A Million Random Digits with 100,000 Normal Deviates. Glencoe, Free Press, 1975.
16. Bloomfield, P. Spectrum analysis of epidemiological data. Paper presented at Fourth Symposium on Statistics and the Environment, Washington, D.C., 1976.
17. Bloomfield, P. Fourier Analysis of Time Series: An Introduction. Wiley, New York, 1976.
18. Box, G. E. P., and Jenkins, G. M. Time Series Analysis Forecasting and Control. Holden-Day, San Francisco, 1970.
19. Box, G. E. P., and Tiao, G. C. Intervention analysis and applications to economic and environmental problems. J. Amer. Statist. Assoc. 70: 70 (1975).
20. Tiao, G. C., Box, G. E. P., and Hamming, W. J. Analysis of Los Angeles photochemical smog data: a statistical overview. J. Air Pollut. Control Assoc. 25: 260 (1975).
21. Tiao, G. C., Box, G. E. P., and Hamming. W. J. A statistical analysis of the Los Angeles ambient carbon monoxide data 1955-1972. J. Air Pollut. Control Assoc., 25: 1129 (1975).

22. Rabiner, L. R., Sambur, M. R., and Schmidt, C. E. Applications of a nonlinear smoothing algorithm to speech processing. IEEE Trans. Acoustics, Speech Signal Proc. 23: No. 6, 552 (1975).

23. Velleman, P. F. Robust Non-Linear Data Smoothing. Tech. Report No. 89, Series 2, Dept. of Statistics, Princeton Univ., 1975.

24. Beaton, E., and Tukey, J. W. The fitting of power series, meaning polynomials, illustrated on band-spectroscopic data. Technometrics, 16: 147 (1974).

25. Hocking, R. R. The analysis and selection of variables in linear regression. Biometrics 32: 1 (1976).

26. Gnanadesikan, R., and Kettenring, J. R. Robust estimates, residuals and outlier detection with multiresponse data. Biometrics 28: 81 (1972).

27. Andrews, D F. A robust method for multiple linear regression. Technometrics 16: 523 (1974).

28. Andrews, D. F., et al. Robust estimates of location: Survey and Advances. Princetion Univ. Press, Princeton, N. J., 1972.

29. Daniel, C. and Wood, F. S. Fitting Equations to Data. Wiley, New York, 1971.

30. Diaconis, P. Measuring the effect of the fuel crisis. Paper given at AAAS meeting, Boston, Mass., 1976.

31. Brown, S. M., et al. Effect on mortality of the 1974 fuel crisis. Nature 257: 306 (1975).

32. Schwing, R. C., and McDonald, G. C. Measures of Association of some air pollutants, natural ionizing radiation and cigarette smoking with mortality rates. Paper presented at International Symposium on Recent Advances in the Assessment of the Health Effects of Environmental Pollution, Paris, France, 1974.

33. McDonald, G. C., and Schwing, R. C. Instabilities of regression estimates relating air pollution to mortality. Technometrics 11: 763 (1973).

34. Barlow, R. E., et al. Statistical Inference under Order Restrictions. Wiley, New York, 1972.

35. Gallant, A. R. Nonlinear regression. Amer. Statistician 29: 73 (1973).

36. Anscombe, F. J., Graphs in statistical analysis. Amer. Statistician 27: 17 (1973).

37. Wilk, M. B., and Gnanadeskan, R. Probability plotting methods for the analysis of data. Biometrika 55: 1 (1968).

38. Cleveland, W. S., and Kleimer, B. Enhancing scatter plots with curves of moving statistics. Technometrics 17: 447 (1975).

39. McNeil, D. R. Statistical models and stochastic models. Proceedings of a SIMS Conference on Epidemiology, Alta, Utah, 1974.

40. Nelson, W. C., Hasselblad, V., and Lowrimond, G. R. Statistical aspects of a community health and environmental surveillance system. VIth Berkeley Symposium, Vol. 6, 1972, p. 125.

41. NIEHS. Man's health and the environment—some research needs. First Task Force on Research Planning in Environmental Health Science. U. S. DHEW, Washington, D. C., 1970.

42. Pratt, J. W. Statistical and Mathematical Aspects of Pollution Problems. (Statistics Textbooks and Monographs, Vol. 6), Dekker, New York, 1974, p. 384.